

Optimizing ROC Curves with a Sort-Based Surrogate Loss for Binary Classification and Changepoint Detection

Jonathan Hillman

*School of Informatics, Computing, and Cyber Systems
Northern Arizona University
Flagstaff, AZ, USA*

JDH553@NAU.EDU

Toby Dylan Hocking

*School of Informatics, Computing, and Cyber Systems
Northern Arizona University
Flagstaff, AZ, USA*

TOBY.HOCKING@NAU.EDU

Editor: Animashree Anandkumar

Abstract

Receiver Operating Characteristic (ROC) curves are useful for evaluating binary classification models, but difficult to use for learning since the Area Under the Curve (AUC) is a piecewise constant function of predicted values. ROC curves can also be used in other problems with false positive and true positive rates such as changepoint detection. We show that in this more general context, the ROC curve can have loops, points with highly sub-optimal error rates, and AUC greater than one. This observation motivates a new optimization objective: rather than maximizing the AUC, we would like a monotonic ROC curve with $AUC=1$ that avoids points with large values for $\text{Min}(FP, FN)$. We propose an L1 relaxation of this objective that results in a new surrogate loss function called the AUM, short for Area Under $\text{Min}(FP, FN)$. Whereas previous loss functions are based on summing over all labeled examples or pairs, the AUM requires a sort and a sum over the sequence of points on the ROC curve. We show that AUM directional derivatives can be efficiently computed and used in a gradient descent learning algorithm. In our empirical study of supervised binary classification and changepoint detection problems, we show that our new AUM minimization learning algorithm results in improved AUC and speed relative to previous baselines.

Keywords: AUC, ROC, loss functions, gradient descent, optimization

1. Introduction

In supervised machine learning problems such as binary classification and changepoint detection, the goal is to learn a function for accurately predicting presence or absence of a class label. In binary classification there is a prediction for each example; in changepoint detection there is a prediction (change or not) in between each data point in a sequence. There are numerous ways to analyze prediction performance, but the simplest way is to calculate accuracy, which is the number or proportion of correctly classified labels. However using accuracy as the evaluation metric can be problematic for data sets with imbalanced labels or for which the desired weighting of the labels is unknown (Cortes and Mohri, 2004; Menon et al., 2013).

A popular approach for comparing models in this context is by analyzing their Receiver Operating Characteristic (ROC) curves, which are plots of True Positive Rate (TPR) versus False Positive Rate (FPR) that have long been used in the signal processing literature (Egan and Egan, 1975). For data with n labeled examples, most algorithms compute a predicted value $\hat{y}_i \in \mathbb{R}$ for each labeled example $i \in \{1, \dots, n\}$. In binary classification the predicted value \hat{y}_i is then compared to the threshold of zero to classify the example as either positive or negative. True positives are examples i with positive labels and positive predictions, whereas false positives are examples i with negative labels

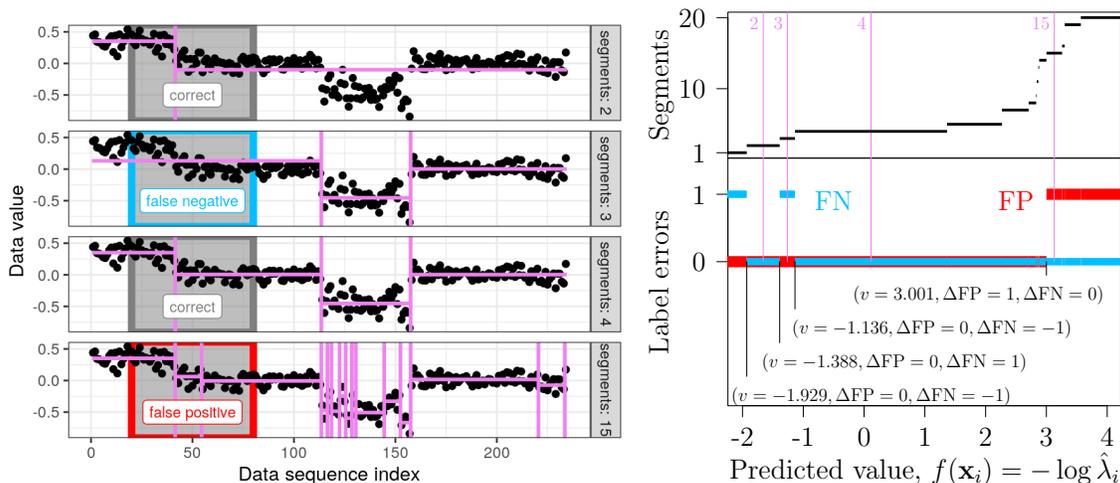


Figure 1: Example labeled changepoint detection data for which the FN function is non-monotonic. **Left:** a data sequence (black dots) with one label (grey rectangle) in which there should be exactly one predicted changepoint (false negative for no changes when segments=3, false positive for two changes when segments=15). **Right:** selected number of segments (top) and number of label errors (bottom) as a function of predicted values, with vertical violet lines indicating the models shown on the left, and vertical black lines for breakpoints in the piecewise constant error functions (v for predicted value, ΔFP , ΔFN for changes in error functions at that value).

and positive predictions. A vector of predicted values $\hat{\mathbf{y}} = [\hat{y}_1 \cdots \hat{y}_n]^\top \in \mathbb{R}^n$, one element for each labeled example i , can be mapped to a single (FPR, TPR) point in ROC space. Each different point on the ROC curve is obtained by adding a real-valued constant $c \in \mathbb{R}$ to each prediction in that vector, $\hat{\mathbf{y}} + c$. Large constants c result in FPR=TPR=1 and small constants result in FPR=TPR=0. Therefore the Area Under the Curve (AUC) is an evaluation metric which accounts for all possible constants or possibilities for thresholding the prediction vector. In binary classification, a perfect model has an AUC of 1, a constant model has an AUC of 0.5, and the minimum AUC is 0.

When AUC is used as the evaluation metric in machine learning, the best algorithm is typically defined as the one that maximizes test AUC. In such experiments, we would also like to use the AUC as the objective function for training the model (under the assumption that if the train and test sets are similar, maximizing train AUC should result in maximizing test AUC). However, since the AUC is a piecewise constant function of the predicted values, its gradient is zero almost everywhere, and it is therefore impossible to directly optimize using gradient descent algorithms.

In addition to studying binary classification, we also study supervised changepoint detection, in which we can also compute FPR and TPR. For binary classification, we can compute a prediction vector which maximizes AUC in the same way as maximizing accuracy — predict a positive value for each positive label, and a negative value for each negative label. For changepoint detection, it is also trivial to compute a prediction vector that maximizes accuracy, but it can be non-trivial to compute a prediction vector that maximizes AUC. The previous observations motivate this paper, which explores a new loss function and corresponding learning algorithm which we empirically show results in AUC maximization.

1.1 Contributions and organization

Our main contribution is the AUM, which is a new surrogate loss function defined as the Area Under the Minimum of false positives and false negatives as a function of the prediction threshold. Whereas previously proposed loss functions for binary classification can be interpreted as convex surrogates of the zero-one loss (summed over examples) or Mann-Whitney statistic (summed over pairs of examples), our proposed AUM loss is an L1 relaxation of the total $\text{Min}(\text{FP}, \text{FN})$ (summed over points on the ROC curve, or distinct intervals of prediction thresholds).

In Section 2 we give precise definitions of the AUC and our new AUM loss function. In Section 3 we give an efficient algorithm for computing directional derivatives of the AUM with respect to predicted values, which we propose using in gradient descent learning algorithms. In Section 4 we provide an empirical study of supervised binary classification and changepoint detection problems, showing that minimizing the AUM corresponds to optimizing the points in the ROC curve. (with respect to train and test sets). Section 5 concludes with a discussion of the significance and novelty of our findings.

1.2 Related work

Evaluating learned binary classification models. ROC curve analysis is a classical evaluation technique with origins in the signal processing literature (Egan and Egan, 1975). Provost and Fawcett (1997) introduced a ROC convex hull method for evaluation. Ling et al. (2003) provide a proof of AUC consistency, whereas Hand (2009) showed that it is incoherent with respect to classification costs, and proposed the “H measure” as an alternative. Martínez-Cambolor et al. (2017) proposed a generalization to predicted values which have a non-monotone relationship with the desired class label.

Margin losses for binary classification. Margin losses can be interpreted as weighted convex surrogates of the zero-one loss, summed over all examples. Ferri et al. (2002); Cortes and Mohri (2004) proposed expressions for the expected value and the variance of the AUC for a fixed error rate, in an attempt to maximize the AUC with different nonlinear algorithms. Wang et al. (2015) incorporated unlabeled data to make an unsupervised AUC maximization algorithm. Other algorithms have been proposed to obtain approximations of the global AUC value (Rakotomamonjy, 2004; Herschtal and Raskutti, 2004; Herschtal et al., 2006; Calders and Jaroszewicz, 2007). Han and Zhao (2010) proposed an active learning algorithm for computing a linear model that maximizes the AUC. Zhao et al. (2011) implemented an online learning algorithm that maximizes the AUC. Menon et al. (2013) analyzed statistical consistency under class imbalance, using a mean of true positive and true negative rates. Scott (2012) studied calibration for margin-based losses for binary classification.

Pairwise losses for binary classification. Bamber (1975) is credited as the first to prove the equivalence of the ROC-AUC and the Mann-Whitney test statistic (Mann and Whitney, 1947). This equivalence has led many authors to propose algorithms based on loss functions that are convex surrogates of the Mann-Whitney statistic. Yan et al. (2003) proposed a global approximation of AUC which was then used in several other algorithms. For example, Castro and Braga (2008) used that approximation to propose the AUCtron algorithm which learns a linear model. Joachims (2005) proposed a quadratic time support vector machine algorithm for AUC maximization based on a pairwise loss function, and Freund et al. (2003) proposed a similar approach based on boosting. Narasimhan and Agarwal (2013) extended this approach to the partial AUC, which is defined as the area under the curve between two false positive rates (not necessarily 0 and 1). Kotlowski et al. (2011) analyze how much risk and regret is increased for balanced margin-based losses, compared with pairwise losses. Rudin et al. (2005) showed that in some cases margin-based losses yield the same solution as a pairwise loss. For stochastic optimization of the AUC, Ying et al. (2016) proposed solving a saddle point problem involving a pairwise square loss, and Yuan et al. (2020) proposed to extend this approach to the squared hinge loss.

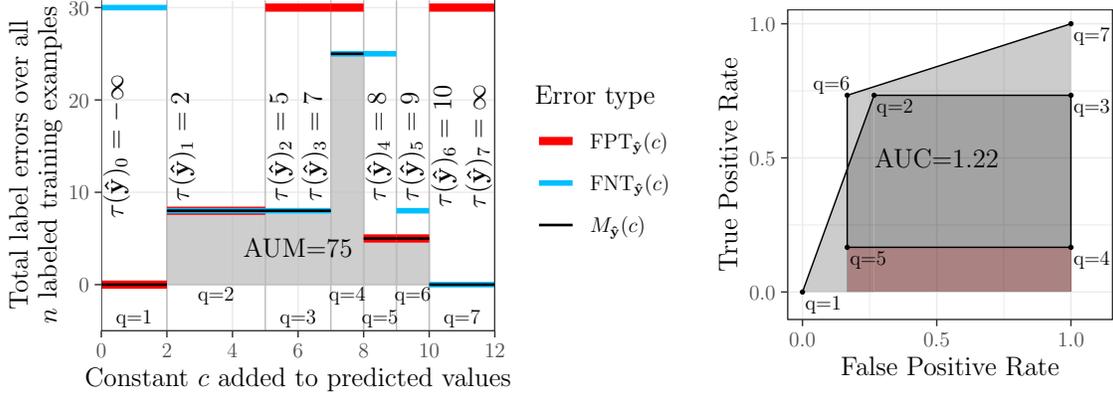


Figure 2: Synthetic example showing how non-monotonic total FP/FN functions (left) can result in a looping ROC curve (right). The q values are indices in the sequence of points on the ROC curve (and corresponding intervals of the FP/FN functions). **Left:** the AUM (grey area) is defined as the area under the minimum of total FP and FN functions, $M_{\hat{y}}(c) = \min\{FPT_{\hat{y}}(c), FNT_{\hat{y}}(c)\}$. For each interval $q \in \{1, \dots, 7\}$ of constant total error, the $\tau(\hat{y})_q$ is the largest threshold in that interval. **Right:** $AUC \geq 1$ due to the loop which results in double counting the dark grey area (but single counting the red area which is positive counted twice and negative counted once).

Changepoint detection. There are many algorithms for unsupervised changepoint detection (Aminikhanghahi and Cook, 2017; van den Burg and Williams, 2020). In supervised changepoint detection, learning is often limited to grid search (Hocking et al., 2013, 2016; Liehrmann et al., 2021). More sophisticated learning algorithms use linear models and non-linear decision trees that minimize convex surrogates of the label error (Rigaill et al., 2013; Hocking et al., 2014, 2015; Drouin et al., 2017; Hocking and Bourque, 2020). ROC curves and AUC are used to evaluate prediction accuracy of learned penalty functions (Maidstone et al., 2016; Hocking et al., 2020).

2. Models and Definitions

We begin by reviewing supervised binary classification and changepoint detection, then give definitions for AUC and the new AUM loss function.

2.1 Review of supervised binary classification

In supervised binary classification we are given a set of n labeled training examples, $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ where $\mathbf{x}_i \in \mathbb{R}^p$ is an input feature vector for one example and $y_i \in \{-1, 1\}$ is a binary output/label. The goal of binary classification is to learn a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ which is used to compute real-valued predictions $\hat{y}_i = f(\mathbf{x}_i)$ with the same sign as the corresponding label y_i . Margin-based learning algorithms such as logistic regression and support vector machines involve minimizing a convex surrogate loss function $\ell : \mathbb{R} \rightarrow \mathbb{R}$, summed over all training examples:

$$\mathcal{L}(f) = \sum_{i=1}^n \ell[y_i f(\mathbf{x}_i)]. \tag{1}$$

Large margin values $y_i f(\mathbf{x}_i) > 0$ yield correct predictions, whereas small values $y_i f(\mathbf{x}_i) < 0$ yield incorrect predictions. Pairwise loss functions involve minimizing a convex loss summed over all pairs

of positive and negative examples,

$$\mathcal{L}(f) = \sum_{i:y_i=-1} \sum_{j:y_j=1} \ell[f(\mathbf{x}_j) - f(\mathbf{x}_i)]. \quad (2)$$

Large pairwise difference values $f(\mathbf{x}_j) - f(\mathbf{x}_i) > 0$ yield correctly ranked pairs, whereas small pairwise difference values $f(\mathbf{x}_j) - f(\mathbf{x}_i) < 0$ yield incorrectly ranked pairs. Typical choices for the convex surrogate loss function $\ell(z)$ include logistic $\log[1 + \exp(-z)]$, linear hinge $(1 - z)_+$, and squared hinge $(1 - z)_+^2$, where $(\cdot)_+$ is the positive part function. After a function f has been learned using the training data, it can be evaluated by computing non-convex evaluation metrics such as the zero-one loss or the AUC with respect to a held-out test set.

2.2 Review of supervised changepoint detection

Supervised changepoint detection is a problem that occurs in various fields of study such as genomics (Hocking et al., 2014) neuroscience (Jewell et al., 2019), and electrocardiogram data analysis (Fotouhinasab et al., 2021). The goal is to accurately detect changepoints which are consistent with the partial labels that are given in a set of data sequences measured over space or time. Here we give a brief overview of supervised changepoint detection; for details see (Rigaill et al., 2013; Truong et al., 2017). We assume for each labeled training example $i \in \{1, \dots, n\}$ there is a corresponding data sequence vector \mathbf{z}_i and label set L_i . For example in Figure 1 we show a data sequence \mathbf{z}_i with a label set L_i that contains one positive label (grey region in which there should be exactly one predicted changepoint). Dynamic programming algorithms can be used on the data sequence \mathbf{z}_i to compute a path of optimal changepoint models $\hat{\mathbf{m}}_i^k$ for different model sizes $k \in \{1, 2, \dots\}$ (Maidstone et al., 2016). For example in Figure 1 (left) we show four models in the path (with $k = 2, 3, 4, 15$ segments). The label set L_i can be used to compute the number of false positive and false negative labels with respect to any predicted set of changepoints (false positives for too many changepoints, false negatives for not enough changepoints). Each example i also has a model selection function $k_i^* : \mathbb{R}_0^+ \rightarrow \{1, 2, \dots\}$ which maps a non-negative penalty value $\hat{\lambda}_i$ to a selected model size $k_i^*(\hat{\lambda}_i)$ (Figure 1, right bottom). We assume there is a fixed feature map ϕ which can be used to compute a feature vector $\mathbf{x}_i = \phi(\mathbf{z}_i) \in \mathbb{R}^p$ for each labeled example (typical features include measures of data sequence size, estimates of variance/noise, etc).

We want to learn a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ which inputs a feature vector and outputs a real-valued prediction that is used as a negative log penalty value, $f(\mathbf{x}_i) = -\log \hat{\lambda}_i$. The goal is to predict model sizes $k_i^*(\hat{\lambda}_i)$ that result in minimal label errors. Since the label error function is non-convex like the zero-one loss in binary classification (Figure 1, right top), previous learning algorithms instead use the gradient with respect to a convex surrogate such as a hinge loss (Rigaill et al., 2013; Drouin et al., 2017) or a censored regression loss (Barnwal et al., 2020).

2.3 Definition of false positive and negative functions

In this paper, we assume the following general learning context in which supervised binary classification and changepoint detection are specific examples. For each labeled training example i , we have one or more labels such that there are at most $\text{FNP}_i \in \mathbb{Z}_+ = \{0, 1, \dots\}$ false negatives possible and $\text{FPP}_i \in \mathbb{Z}_+$ false positives possible. Given a real-valued prediction $\hat{y}_i = f(\mathbf{x}_i) \in \mathbb{R}$, we can use the labels to compute the number of predicted false positives $\text{FP}_i(\hat{y}_i) \in \{0, \dots, \text{FPP}_i\}$ and false negatives $\text{FN}_i(\hat{y}_i) \in \{0, \dots, \text{FNP}_i\}$. The FP_i, FN_i functions return the number of false positive and false negative labels for a given predicted value \hat{y}_i .

Exact representation using breakpoints. By convention we assume that the FP_i, FN_i functions are piecewise constant and right continuous; that the FP_i functions start at zero, $\lim_{x \rightarrow -\infty} \text{FP}_i(x) = 0$; and the FN_i functions end at zero, $\lim_{x \rightarrow \infty} \text{FN}_i(x) = 0$. These assumptions ensure that (1) our proposed AUM loss function will always be finite, and (2) it can be computed efficiently by representing

the error functions exactly using a finite set of breakpoints. For each breakpoint tuple $(v, \Delta\text{FP}, \Delta\text{FN})$, $v \in \mathbb{R}$ is a predicted value threshold where there are changes $\Delta\text{FP}, \Delta\text{FN}$ (discontinuity) in the error functions, $\text{FP}_i(v) - \lim_{x \rightarrow v^-} \text{FP}_i(x) = \Delta\text{FP}$ and similar for FN. For example, in Figure 1 the error functions can be exactly represented by a set of four such breakpoints. These breakpoints will be used when computing our proposed loss function and learning algorithm.

Case of binary classification. In the case of binary classification, for all positive examples $i : y_i = 1$ we have $\text{FPP}_i = 0$, $\text{FNP}_i = 1$, $\text{FP}_i(\hat{y}) = 0$, $\text{FN}_i(\hat{y}) = I(\hat{y} < 0)$, where I is the indicator function (outputs 1 if argument is true, 0 otherwise). For all negative examples $i : y_i = -1$, $\text{FPP}_i = 1$, $\text{FNP}_i = 0$, $\text{FP}_i(\hat{y}) = I(\hat{y} \geq 0)$, $\text{FN}_i(\hat{y}) = 0$. Note that FP_i is either constant/zero (for positive examples) or non-decreasing (for negative examples), and FN_i is either constant/zero (for negative examples) or non-increasing (for positive examples). Since the prediction threshold is always zero in binary classification, these functions can be exactly represented by the breakpoint $(v = 0, \Delta\text{FP} = 0, \Delta\text{FN} = -1)$ for all positive examples, and $(v = 0, \Delta\text{FP} = 1, \Delta\text{FN} = 0)$ for all negative examples.

Case of changepoint detection. In changepoint detection, we have more general FP_i and FN_i functions that can be non-monotonic, with data-dependent thresholds v that can be computed in advance of learning f . For example, in Figure 1, we show a data sequence with one positive label (in which there should be exactly one predicted changepoint). Predicting no changepoint in this label results in a false negative, and predicting two changepoints in this label results in a false positive. Therefore, we have $\text{FPP}_i = \text{FNP}_i = 1$ for this particular example i ; the false positive function is non-decreasing, $\text{FP}_i(\hat{y}_i) \approx I(\hat{y}_i \geq 3.001)$, and the false negative function is not monotonic, $\text{FN}_i(\hat{y}_i) \approx I[\hat{y}_i \in (\infty, -1.929) \cup (-1.388, -1.136)]$. When the false negative function is non-monotonic, the true positive rate is also non-monotonic (the ROC curve can move down as well as up when prediction constant is increased). Given a pre-computed path of changepoint models with loss/size values, the exact breakpoints $(v, \Delta\text{FP}, \Delta\text{FN})$ in such error functions can be efficiently computed using a linear time algorithm (Hocking and Vargovich, 2020).

2.4 Definition of ROC curve and AUC

In this section, we show how the previously defined functions can be used to define the ROC curve and the AUC. Given a vector of predictions $\hat{\mathbf{y}} \in \mathbb{R}^n$, and a constant $c \in \mathbb{R}$ added to that vector, we define the False Positive Total (FPT) and False Negative Total (FNT) functions as

$$\text{FPT}_{\hat{\mathbf{y}}}(c) = \sum_{i=1}^n \text{FP}_i(\hat{y}_i + c), \quad (3)$$

$$\text{FNT}_{\hat{\mathbf{y}}}(c) = \sum_{i=1}^n \text{FN}_i(\hat{y}_i + c). \quad (4)$$

The corresponding False Positive Rate (FPR) and True Positive Rate (TPR) functions are

$$\text{FPR}_{\hat{\mathbf{y}}}(c) = \frac{1}{\sum_{i=1}^n \text{FPP}_i} \text{FPT}_{\hat{\mathbf{y}}}(c), \quad (5)$$

$$\text{TPR}_{\hat{\mathbf{y}}}(c) = 1 - \frac{1}{\sum_{i=1}^n \text{FNP}_i} \text{FNT}_{\hat{\mathbf{y}}}(c). \quad (6)$$

The ROC curve for a given prediction vector $\hat{\mathbf{y}}$ is the plot of $\text{TPR}_{\hat{\mathbf{y}}}(c)$ versus $\text{FPR}_{\hat{\mathbf{y}}}(c)$ as the constant c is varied from $-\infty$ to ∞ . The AUC is the integral,

$$\text{AUC}(\hat{\mathbf{y}}) = \int \text{TPR}_{\hat{\mathbf{y}}}(c) d\text{FPR}_{\hat{\mathbf{y}}}(c). \quad (7)$$

Assuming the $FPT_{\hat{y}}(c)$ and $FNT_{\hat{y}}(c)$ functions are piecewise constant (as in binary classification and changepoint detection) then the ROC curve can be described as a sequence of Q points $\{(fpr(\hat{y})_q, tpr(\hat{y})_q)\}_{q=1}^Q$ in ROC space (note that this sequence of points depends on the predicted values \hat{y}). The first point $q = 1$ corresponds to prediction threshold $t = -\infty$ which results in $TPR=0$ and $FPR=0$; the last point $q = Q$ is for $t = \infty$ which results in $TPR=1$ and $FPR=1$ (Figure 2). The AUC can be computed using this sequence,

$$AUC(\hat{y}) = \sum_{q=2}^Q (fpr(\hat{y})_q - fpr(\hat{y})_{q-1})(tpr(\hat{y})_{q-1} + tpr(\hat{y})_q)/2. \quad (8)$$

In binary classification we have $fpr(\hat{y})_{q-1} \leq fpr(\hat{y})_q$ which means while tracing the ROC curve from $FPR=TPR=0$ to $FPR=TPR=1$ there are no moves to the left, and all the terms in the sum above are positive. In fact this is also true of changepoint detection problems for which all the FP_i functions are non-decreasing. In these cases the ROC curve is monotonic, with $AUC(\hat{y}) \in [0, 1]$ for any predicted values \hat{y} .

2.5 Looping curves motivate optimizing ROC points in changepoint problems

In the context of changepoint detection, it is possible to have non-monotonic ROC curves with loops, when there are data sequences with non-hierarchical changepoint models. For example, Figure 1 (left) shows a data sequence with a non-hierarchical changepoint model, because the changepoint which was detected for a model size of two segments is not present at the model size of three segments. Because there is a positive label for that changepoint, the FN_i function is non-monotonic ($FN=1$ for model sizes of 1 and 3 segments, $FN=0$ for model sizes of 2 and 4 segments). In this case while tracing the ROC curve from $FPR=TPR=0$ to $FPR=TPR=1$, there are moves up as usual, but also down. Similarly, if some of the FP_i functions are not monotonically increasing, then it is possible to have $fpr(\hat{y})_{q-1} > fpr(\hat{y})_q$ which results in a move to the left, and a negative term in the AUC equation (8). If there are non-monotonic FN_i and FP_i functions, then it is possible to have loops in the ROC curve, and the AUC is therefore not necessarily bounded between 0 and 1.

For a simple synthetic example, consider the FP/FN functions and corresponding ROC curves in Figure 2. In this example, the ROC curve has area ≈ 1.22 because it contains a loop which double-counts a large portion of ROC space. This looping implies existence of ROC points q which are highly sub-optimal in terms of TPR and/or FPR (for example, $q \in \{2, 3, 4, 5\}$ in Figure 2). By “sub-optimal” we mean ROC points that occur more toward the bottom right of ROC space, where both FP and FN are large, as is $\min\{FP, FN\}$. Thus we observe an association between looping ROC curves, AUC greater than one, and existence of some thresholds or points on the ROC curve with highly sub-optimal error rates.

Although this seems like a rare phenomenon, we have seen this occur in real data. For example, in Figure 3 (left) we show error curves for two labeled examples from another real changepoint detection data set. One example has a positive label that results in a non-monotonic false negative function, and the other example has a negative label that results in a non-monotonic false positive function. When the predicted value for the negative example is about 5 greater than the predicted value for the positive example, we observe $AUC=2$ (Figure 3, right). This is the result of a loop in ROC space, indicating that there is some prediction threshold that results in $(FPR=1, TPR=0)$ which corresponds to 100% error and 0% accuracy. Therefore, unconstrained maximization of the AUC may not be a desirable optimization objective, because it can result in ROC curves with loops and points in the lower right, with large $\min\{FP, FN\}$ values. Instead, we would prefer to maximize the AUC subject to the constraint that the ROC curve is monotonic, so that there are no highly sub-optimal ROC points in the lower right. We therefore propose as an optimization objective to minimize the total $\min\{FP, FN\}$ over all ROC points, which we interpret as attempting to maximize AUC subject to a monotonicity constraint on the sequence of ROC points. This objective means

that ideal ROC points would be moved away from the lower right, toward the upper left, with the best case being a monotonic ROC curve with AUC=1. We formalize this idea in the next section, and provide a new surrogate loss function for this objective.

2.6 Proposed surrogate loss function (AUM)

In this section we propose the AUM loss function, which is short for Area Under Min(FP, FN). The intuition behind the AUM is that we want to minimize the number of prediction thresholds that result in large error rates. To formally define the AUM we must first define the minimum of total false positives and false negatives for a given constant c added to predicted values,

$$M_{\hat{y}}(c) = \min\{\text{FPT}_{\hat{y}}(c), \text{FNT}_{\hat{y}}(c)\}. \quad (9)$$

Note that the definition above uses total counts of false positives and false negatives; this corresponds to the *AUM.count* method we study in the results section. Another variant, which we call *AUM.rate* in the results, defines the min using normalized rates (between 0 and 1) rather than absolute counts,

$$M_{\hat{y}}(c) = \min\{\text{FPR}_{\hat{y}}(c), \text{FNR}_{\hat{y}}(c)\}. \quad (10)$$

Using either of the two definitions of the min above, we define the AUM as

$$\text{AUM}(\hat{y}) = \int_{-\infty}^{\infty} M_{\hat{y}}(c) dc. \quad (11)$$

The definition (11) says that the AUM is computed by integrating the min function $M_{\hat{y}}(c)$ over all possible constants $c \in \mathbb{R}$ that could be added to predicted values. Geometrically, this corresponds to the area under the minimum of total false positive and false negative functions (Figure 2, left). Note that the AUM is well-defined and can be computed for any FP_i, FN_i functions (even those which are non-monotonic, which can occur with non-hierarchical changepoint models).

2.7 Interpretation as L1 relaxation of total Min(FP, FN) over ROC points

Computing the AUM is similar to the AUC, in that we must compute error/accuracy rates for each possible constant c which could be added to predicted values. First let $\{(\text{fpt}(\hat{y})_q, \text{fnt}(\hat{y})_q, \tau(\hat{y})_q)\}_{q=1}^Q$ be a sequence of Q tuples, each of which corresponds to a point on the ROC curve (Figure 2, right). The fpt/fnt are false positive/negative totals whereas τ are values such that there is a change/threshold at $M_{\hat{y}}(\tau)$. As shown in Figure 2 we assume these values are increasing, $-\infty = \tau(\hat{y})_0 < \dots < \tau(\hat{y})_Q = \infty$. For each $q \in \{1, \dots, Q\}$ there is a corresponding interval of values c between $\tau(\hat{y})_{q-1}$ and $\tau(\hat{y})_q$ such that $\text{FPT}_{\hat{y}}(c) = \text{fpt}(\hat{y})_q$ and $\text{FNT}_{\hat{y}}(c) = \text{fnt}(\hat{y})_q$ for all $c \in (\tau(\hat{y})_{q-1}, \tau(\hat{y})_q)$ (Figure 2, left). Then we define the minimum for interval/ROC point q as $m(\hat{y})_q = \min\{\text{fpt}(\hat{y})_q, \text{fnt}(\hat{y})_q\}$. Because the first and last minima are zero, $m(\hat{y})_1 = m(\hat{y})_Q = 0$, the area under the first and last intervals is zero, and the AUM can be computed by summing over all of the other intervals,

$$\text{AUM}(\hat{y}) = \sum_{q=2}^{Q-1} [\tau(\hat{y})_q - \tau(\hat{y})_{q-1}] m(\hat{y})_q. \quad (12)$$

For example, in Figure 2 there are $Q = 7$ tuples/intervals (points on the ROC curve), five of which result in a positive AUM value, resulting in a total AUM of 75.

The AUM can be interpreted as an L1 relaxation of the following non-convex Sum of Min(FP, FN) function,

$$\text{SM}(\hat{y}) = \sum_{q=2}^{Q-1} I[\tau(\hat{y})_q \neq \tau(\hat{y})_{q-1}] m(\hat{y})_q = \sum_{q: \tau(\hat{y})_q \neq \tau(\hat{y})_{q-1}} m(\hat{y})_q. \quad (13)$$

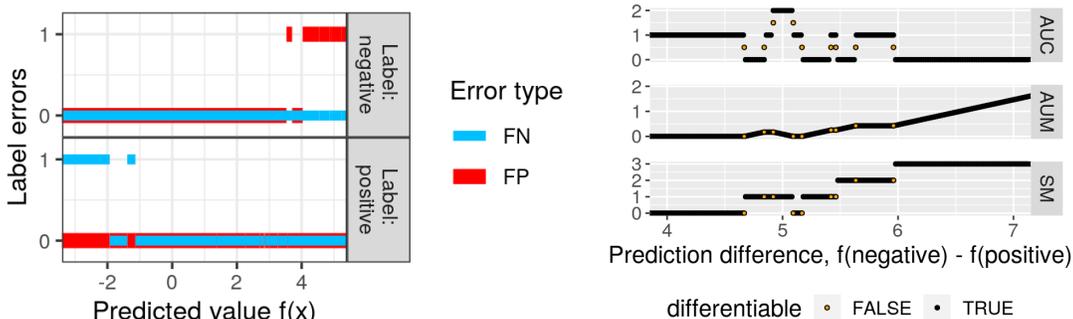


Figure 3: Real data example with AUC greater than one. **Left:** error functions for two labeled examples in a real changepoint detection data set. **Right:** AUC, AUM, and SM values as a function of the difference in predicted values between the two labeled examples. When the difference is around 5, we have $AUC=2$. The AUC and SM are piecewise constant functions of the predicted values, whereas the AUM is a continuous piecewise linear function. At differentiable points, $AUM=0$ implies $AUC=1$ (but the converse is false).

The difference is that the L1 norm of $\tau(\hat{\mathbf{y}})_q - \tau(\hat{\mathbf{y}})_{q-1}$ in (12) has been changed to the L0 pseudo-norm in (13). The indicator function I takes values in zero and one, which means the SM is the sum of $m(\hat{\mathbf{y}})_q$ values over all distinct ROC points q . Geometrically, the SM function measures total distance of distinct points on the ROC curve from the left (FPR=0) or top (TPR=1, FNR=0) of ROC space (whichever is smaller). A visualization of the piecewise constant SM loss as a function of the predicted values \hat{y}_i shows that the AUM is indeed a continuous relaxation that is differentiable almost everywhere (Figure 3, right).

2.8 Properties of AUM

A few interesting properties of the AUM should be immediately apparent.

Finite. Since $FP_i(-\infty) = 0$ and $FN_i(\infty) = 0$ for all i (by assumption), we have $M(\hat{\mathbf{y}} \pm \infty) = 0$, which means the AUM integral (11) is finite, $AUM(\hat{\mathbf{y}}) < \infty$.

Non-negative. Since the label error is never negative, the AUM integral is never negative, $AUM(\hat{\mathbf{y}}) \geq 0$.

Convexity. If all of the FP_i and FN_i functions are monotonic, then the AUM is convex.

Other properties. The AUM is continuous, piecewise linear, and differentiable almost everywhere.

It can be seen that non-differentiable points in the AUM coincide with discontinuities in AUC (Figure 3, right). Furthermore at differentiable points, we see that $AUM(\hat{\mathbf{y}}) = 0$ implies $AUC(\hat{\mathbf{y}}) = 1$. However at non-differentiable points it is possible to have $AUM=0$ with sub-optimal AUC (for example $AUC=0.5$ at the left-most orange dot in Figure 3, right). Additionally there are differentiable points for which $AUC=1$ or 2, and we have $AUM > 0$, which means there are some prediction thresholds which result in $FP=FN=1$ (for example in Figure 3, right, at prediction difference=5 we have $AUM=0.1$ and $AUC=2$). These observations suggest modifying the learning objective from maximizing AUC, to instead minimizing AUM, which results in $AUC=1$ (at differentiable points). We therefore propose an AUM gradient descent learning algorithm in the next section.

3. Algorithms

In this section we propose algorithms for computing the AUM and its gradient, and for learning predictive models using gradient descent. Throughout we use n as the number of examples for which we want to compute predictions and AUM; note that this can be the full data set size (full gradient descent algorithm), or the size of a minibatch (stochastic gradient descent algorithm).

3.1 Details of AUM computation

In previous sections we have assumed that a sequence of Q points on the ROC curve (or equivalently intervals of FP/FN functions) can be computed, and in this section we explain how to do that. Recall from Section 2.3 that the FP_i, FN_i functions have an exact representation in terms of breakpoints. Let there be a total of B breakpoints in the error functions over all n labeled training examples, where each breakpoint $b \in \{1, \dots, B\}$ is represented by the tuple $(v_b, \Delta\text{FP}_b, \Delta\text{FN}_b, \mathcal{I}_b)$. The $\mathcal{I}_b \in \{1, \dots, n\}$ is an example index, so there are changes $\Delta\text{FP}_b, \Delta\text{FN}_b$ at predicted value $v_b \in \mathbb{R}$ in the error functions $\text{FP}_{\mathcal{I}_b}, \text{FN}_{\mathcal{I}_b}$ (Figure 1). For example in binary classification, there are $B = n$ breakpoints (same as the number of labeled training examples); for each breakpoint $b \in \{1, \dots, B\}$ we have $v_b = 0$ and $\mathcal{I}_b = b$. For breakpoints b with positive labels $y_b = 1$ we have $\Delta\text{FP} = 0, \Delta\text{FN} = -1$, and for negative labels $y_b = -1$ we have $\Delta\text{FP} = 1, \Delta\text{FN} = 0$. In changepoint detection we have more general error functions, which may have more than one breakpoint per example. For example the labeled data sequence shown in Figure 1 is a single labeled training example i with error functions that can be represented by four breakpoints b with the same $\mathcal{I}_b = i$ value.

Given a prediction vector $\hat{\mathbf{y}} = [\hat{y}_1 \cdots \hat{y}_n]^\top \in \mathbb{R}^n$ we can compute a prediction threshold $t_b = v_b - \hat{y}_{\mathcal{I}_b}$ for each breakpoint $b \in \{1, \dots, B\}$. The prediction thresholds $t \in \{t_1, \dots, t_B\}$ are where there are changes in the total error functions $\text{FPT}_{\hat{\mathbf{y}}}(t), \text{FNT}_{\hat{\mathbf{y}}}(t)$ defined in equations (3–4). These functions can be exactly represented by the sequences of $b \in \{1, \dots, B\}$ error values

$$\underline{\text{FP}}_b = \sum_{j:t_j < t_b} \Delta\text{FP}_j, \quad (14)$$

$$\overline{\text{FP}}_b = \sum_{j:t_j \leq t_b} \Delta\text{FP}_j, \quad (15)$$

$$\underline{\text{FN}}_b = \sum_{j:t_j \geq t_b} -\Delta\text{FN}_j, \quad (16)$$

$$\overline{\text{FN}}_b = \sum_{j:t_j > t_b} -\Delta\text{FN}_j. \quad (17)$$

The $\underline{\text{FP}}_b, \underline{\text{FN}}_b$ are the error values before the threshold t_b , whereas $\overline{\text{FP}}_b, \overline{\text{FN}}_b$ are the error values after the threshold. We sort the breakpoints by threshold, yielding a permutation $\{s_1, \dots, s_B\}$ of the indices $\{1, \dots, B\}$ such that for every $q \in \{2, \dots, B\}$ we have $t_{s_{q-1}} \leq t_{s_q}$. All of the error values $\underline{\text{FP}}_b, \overline{\text{FP}}_b, \underline{\text{FN}}_b, \overline{\text{FN}}_b$, for every $b \in \{1, \dots, B\}$, can then be computed via a modified cumulative sum (in the forward direction starting with $\underline{\text{FP}}_1 = 0$, and in the backward direction starting with $\overline{\text{FN}}_B = 0$). In terms of the notation of Section 2.7, we have $Q = B + 1$ points (not necessarily unique) on the ROC curve such that $\tau(\hat{\mathbf{y}})_q = t_{s_q}$ and $m(\hat{\mathbf{y}})_q = \min\{\underline{\text{FP}}_{s_q}, \underline{\text{FN}}_{s_q}\} = \min\{\overline{\text{FP}}_{s_{q-1}}, \overline{\text{FN}}_{s_{q-1}}\}$. The AUM can then be computed via

$$\text{AUM}(\hat{\mathbf{y}}) = \sum_{q=2}^B (t_{s_q} - t_{s_{q-1}}) \min\{\overline{\text{FP}}_{s_{q-1}}, \overline{\text{FN}}_{s_{q-1}}\}, \quad (18)$$

$$= \sum_{q=2}^B (t_{s_q} - t_{s_{q-1}}) \min\{\underline{\text{FP}}_{s_q}, \underline{\text{FN}}_{s_q}\}. \quad (19)$$

The equations above state that the AUM can be computed by multiplying each threshold difference $t_{s_q} - t_{s_{q-1}}$ by the minimum below breakpoint s_q (19) or above breakpoint s_{q-1} (18). Since the slowest operation is the sort, the overall time complexity for computing the AUM is $O(B \log B)$.

3.2 Algorithm for computing directional derivatives

First we note that the AUM function is not differentiable everywhere, so the gradient is not defined everywhere (for example, orange dots in Figure 3, right). Second we note that the AUM function can be non-convex, in which case the sub-differential from convex analysis is not defined (Rockafellar, 1970). Instead, we propose an algorithm for computing the AUM directional derivatives, which are defined everywhere. We recall the general definition of a directional derivative.

Definition 1 *Given vectors $\mathbf{x}, \mathbf{v} \in \mathbb{R}^n$ and a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the directional derivative of f at \mathbf{x} in the direction of \mathbf{v} is the function $\nabla_{\mathbf{v}} f : \mathbb{R}^n \rightarrow \mathbb{R}$ given by*

$$\nabla_{\mathbf{v}} f(\mathbf{x}) = \lim_{h \rightarrow 0} \frac{f(\mathbf{x} + h\mathbf{v}) - f(\mathbf{x})}{h}. \quad (20)$$

We would like to compute $\nabla_{\mathbf{v}} \text{AUM}(\hat{\mathbf{y}})$, for a set of directions \mathbf{v} . We are interested in computing the directional derivative along a single dimension $i \in \{1, \dots, n\}$, in either the negative or positive direction, which correspond to using direction vectors \mathbf{v} with -1 or 1 at the i -th position, and zeros at each other position,

$$\mathbf{v}(-1, i) = [0 \ \cdots \ -1 \ \cdots \ 0]^\top, \quad (21)$$

$$\mathbf{v}(1, i) = [0 \ \cdots \ 1 \ \cdots \ 0]^\top. \quad (22)$$

The intuition of these direction vectors is that each will give us the rate of change of AUM, if a single prediction i is either increased or decreased. We propose an algorithm for efficiently computing the following $n \times 2$ matrix of directional derivatives,

$$\mathbf{D}_f(\mathbf{x}) = \begin{bmatrix} \nabla_{\mathbf{v}(-1,1)} f(\mathbf{x}) & \cdots & \nabla_{\mathbf{v}(-1,n)} f(\mathbf{x}) \\ \nabla_{\mathbf{v}(1,1)} f(\mathbf{x}) & \cdots & \nabla_{\mathbf{v}(1,n)} f(\mathbf{x}) \end{bmatrix}^\top. \quad (23)$$

We will compute $\mathbf{D}_{\text{AUM}}(\hat{\mathbf{y}})$, which is the matrix of directional derivatives for a given prediction vector. If we have equality of elements of all rows of this matrix, i.e., $\nabla_{\mathbf{v}(-1,i)} \text{AUM}(\hat{\mathbf{y}}) = \nabla_{\mathbf{v}(1,i)} \text{AUM}(\hat{\mathbf{y}})$ for all i , then the gradient does exist at $\hat{\mathbf{y}}$ and is equal to that value. The following theorem shows how to compute the elements of this matrix.

Theorem 2 *The AUM directional derivatives for a particular example $i \in \{1, \dots, n\}$ can be computed using the following equations.*

$$\nabla_{\mathbf{v}(-1,i)} \text{AUM}(\hat{\mathbf{y}}) = \sum_{b: \mathcal{I}_b = i} \min\{\overline{\text{FP}}_b, \overline{\text{FN}}_b\} - \min\{\overline{\text{FP}}_b - \Delta \text{FP}_b, \overline{\text{FN}}_b - \Delta \text{FN}_b\}, \quad (24)$$

$$\nabla_{\mathbf{v}(1,i)} \text{AUM}(\hat{\mathbf{y}}) = \sum_{b: \mathcal{I}_b = i} \min\{\underline{\text{FP}}_b + \Delta \text{FP}_b, \underline{\text{FN}}_b + \Delta \text{FN}_b\} - \min\{\underline{\text{FP}}_b, \underline{\text{FN}}_b\}. \quad (25)$$

Proof We can compute AUM using either (18) or (19). To compute a directional derivative we need to evaluate $\text{AUM}[\hat{\mathbf{y}} + h\mathbf{v}(d, i)]$. To do that we must use (18) if $d = -1$ and (19) if $d = 1$. First consider the case of $d = 1$, if $t_{s_{q-1}} = t_{s_q}$, and $\mathcal{I}_{s_{q-1}} = i$. In that case the analogous term in $\text{AUM}[\hat{\mathbf{y}} + h\mathbf{v}(d, i)]$ will have $v_{s_{q-1}} - \hat{y}_{s_{q-1}} - h < v_{s_q} - \hat{y}_{s_q}$. Then the min in the corresponding term

Algorithm 1 AUM and Directional Derivatives

- 1: **Input:** Predictions $\hat{\mathbf{y}} \in \mathbb{R}^n$, breakpoints in error functions $v_b, \Delta\text{FP}_b, \Delta\text{FN}_b, \mathcal{I}_b$ for all $b \in \{1, \dots, B\}$.
 - 2: Initialize to zero the AUM $\in \mathbb{R}$ and directional derivative matrix $\mathbf{D} \in \mathbb{R}^{n \times 2}$.
 - 3: $t_b \leftarrow v_b - \hat{y}_{\mathcal{I}_b}$ for all b .
 - 4: $s_1, \dots, s_B \leftarrow \text{SORTEDINDICES}(t_1, \dots, t_B)$.
 - 5: Compute $\underline{\text{FP}}_b, \overline{\text{FP}}_b, \underline{\text{FN}}_b, \overline{\text{FN}}_b$ for all b using s_1, \dots, s_B and (14–17).
 - 6: **for** $b \in \{2, \dots, B\}$ **do**
 - 7: AUM += $(t_{s_b} - t_{s_{b-1}}) \min\{\underline{\text{FP}}_b, \overline{\text{FN}}_b\}$.
 - 8: **for** $b \in \{1, \dots, B\}$ **do**
 - 9: $\mathbf{D}_{\mathcal{I}_b,1} += \min\{\overline{\text{FP}}_b, \overline{\text{FN}}_b\} - \min\{\overline{\text{FP}}_b - \Delta\text{FP}_b, \overline{\text{FN}}_b - \Delta\text{FN}_b\}$
 - 10: $\mathbf{D}_{\mathcal{I}_b,2} += \min\{\underline{\text{FP}}_b + \Delta\text{FP}_b, \underline{\text{FN}}_b + \Delta\text{FN}_b\} - \min\{\underline{\text{FP}}_b, \underline{\text{FN}}_b\}$
 - 11: **Output:** AUM and matrix \mathbf{D} of directional derivatives.
-

must be $\min\{\underline{\text{FP}}_{s_{b-1}} + \Delta\text{FP}_{s_{b-1}}, \underline{\text{FN}}_{s_{b-1}} + \Delta\text{FP}_{s_{b-1}}\}$ and so

$$\begin{aligned} \nabla_{\mathbf{v}(1,i)} \text{AUM}(\hat{\mathbf{y}}) &= \sum_{b=2}^B I[\mathcal{I}_{s_{b-1}} = i] \min\{\underline{\text{FP}}_{s_{b-1}} + \Delta\text{FP}_{s_{b-1}}, \underline{\text{FN}}_{s_{b-1}} + \Delta\text{FP}_{s_{b-1}}\} \\ &\quad - \sum_{b=2}^B I[\mathcal{I}_{s_b} = i] \min\{\underline{\text{FP}}_{s_{b-1}} + \Delta\text{FP}_{s_{b-1}}, \underline{\text{FN}}_{s_{b-1}} + \Delta\text{FP}_{s_{b-1}}\}, \end{aligned} \quad (26)$$

$$\begin{aligned} &= \sum_{b=1}^B I[\mathcal{I}_{s_b} = i] \min\{\underline{\text{FP}}_{s_b} + \Delta\text{FP}_{s_b}, \underline{\text{FN}}_{s_b} + \Delta\text{FP}_{s_b}\} \\ &\quad - \sum_{b=1}^B I[\mathcal{I}_{s_b} = i] \min\{\underline{\text{FP}}_{s_b}, \underline{\text{FN}}_{s_b}\}. \end{aligned} \quad (27)$$

The first equality comes from the definition of the directional derivative and the second re-writes some of the $b - 1$ indices as b . The sum can be extended to start at $b = 1$ since $\underline{\text{FP}}_{s_1} = 0$ (the first min is zero). Finally, re-writing the sums and removing the indicator functions obtains the desired result. The proof for the case of $d = -1$ is analogous. \blacksquare

3.3 Pseudocode and complexity analysis

We propose to compute the matrix of directional derivatives using Algorithm 1 which inputs a prediction vector $\hat{\mathbf{y}} \in \mathbb{R}^n$ and an exact description of the error functions in terms of breakpoints $v_b, \Delta\text{FP}_b, \Delta\text{FN}_b, \mathcal{I}_b$. The first step is to compute thresholds $t_b = v_b - \hat{y}_{\mathcal{I}_b}$ in the total error functions (line 3) which are then sorted (line 4). The sorted indices are then used to compute modified cumulative sums of false positives and false negatives, $\underline{\text{FP}}_b, \overline{\text{FP}}_b, \underline{\text{FN}}_b, \overline{\text{FN}}_b$, for each breakpoint b (line 5). Note that for FP we start with $\underline{\text{FP}}_1 = 0$ and add in the forward direction, whereas for FN we start with $\overline{\text{FN}}_B = 0$ and subtract in the reverse direction. Each iteration of the for loop over intervals of threshold values (line 6) adds to the AUM (line 7). Each iteration of the for loop over breakpoints b (line 8) adds to row \mathcal{I}_b of the directional derivative matrix. The first column stores the derivative in the negative direction (line 9) and the second column stores the derivative in the positive direction (line 10). The time complexity of Algorithm 1 is log-linear $O(B \log B)$ because of the sort (line 4). Since there is at least one breakpoint in each of the n example-specific error functions, the total number of breakpoints $B \geq n$. Therefore, computing the AUM and its directional

derivatives is $O(n \log n)$ time, asymptotically slower than $O(n)$ margin-based loss functions by only a small log factor, and substantially faster than $O(n^2)$ pairwise loss functions.

3.4 Gradient descent algorithm for predicted values

We propose gradient descent optimization algorithms that use the AUM directional derivatives computed using Theorem 2 and Algorithm 1. First to study how minimizing the train AUM affects the train AUC, we propose to optimize the n -vector of predictions $\hat{\mathbf{y}}$. When the AUM is non-convex (in the case of changepoint problems with non-monotonic FP_i and FN_i functions), the initialization of the algorithm is important, because the algorithm will converge to a local (not global) minimum. Therefore when optimizing the vector of predictions, we propose initializing each \hat{y}_i to a value with minimum label errors,

$$\hat{y}_i^{(0)} = \arg \min_x \text{FP}_i(x) + \text{FN}_i(x). \quad (28)$$

After initialization we need to compute a descent direction; typically the negative gradient is used. Recall that since the AUM has non-differentiable points, the gradient is not defined at these points. The “gradient” we propose to use is the mean of the two columns of the directional derivative matrix, $\bar{\mathbf{D}}_{\text{AUM}}(\hat{\mathbf{y}}) \in \mathbb{R}^n$, with each element $i \in \{1, \dots, n\}$ of this vector defined as $[\nabla_{\mathbf{v}(-1,i)}\text{AUM}(\hat{\mathbf{y}}) + \nabla_{\mathbf{v}(1,i)}\text{AUM}(\hat{\mathbf{y}})]/2$. When running our gradient descent algorithm, we have observed that it empirically almost always stays at differentiable points (that is, columns of directional derivative matrix are equal). However, we have observed a few cases where the gradient descent algorithm visits non-differentiable points, for which there is at least one row in the directional derivative matrix with entries that are not equal. Examples of such directional derivative rows $[\nabla_{\mathbf{v}(-1,i)}\text{AUM}(\hat{\mathbf{y}}), \nabla_{\mathbf{v}(1,i)}\text{AUM}(\hat{\mathbf{y}})]$ that we have observed include $[-0.0019, -0.0011]$ and $[-0.0006, 0]$, both of which indicate that the loss would increase if the predicted value is decreased. We also perform line search via grid search in order to obtain a step size $\alpha^{(j)}$ which results in the largest decrease in AUM. Also let $\beta^{(j)}$ be an intercept or threshold with minimal error after taking the line search step (it only affects the label error/accuracy and not the AUM). Note that this intercept can be efficiently computed at the same time as the AUM and its directional derivative matrix, by a simple linear scan over all thresholds. We then perform gradient descent updates for each iteration $j \in \{0, 1, \dots\}$ via

$$\hat{\mathbf{y}}^{(j+1)} = \hat{\mathbf{y}}^{(j)} - \alpha^{(j)} \bar{\mathbf{D}}_{\text{AUM}}(\hat{\mathbf{y}}^{(j)}) + \beta^{(j)}. \quad (29)$$

3.5 Gradient descent learning algorithms for making predictions

In the context of making predictions on a held-out test set using a linear model $f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x} + \beta$, the parameters to learn are a weight vector $\mathbf{w} \in \mathbb{R}^p$ to optimize via gradient descent, and an intercept $\beta \in \mathbb{R}$ which can be optimized via a linear scan to find a threshold with minimum label error. Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ be the feature/input matrix, so $\mathbf{X}\mathbf{w} + \beta \in \mathbb{R}^n$ is the vector of predicted values on the train set. In the context of binary classification problems we use an initialization $\mathbf{w}^{(0)}$ near zero. In the context of changepoint detection problems we consider an initialization $\mathbf{w}^{(0)}$ based on minimizing a convex squared hinge loss with L1 regularization (Rigail et al., 2013). This convex loss function has a minimum for each example at predicted values that achieve minimum label errors, so we expect this initialization to have large accuracy but not necessarily large AUC. Again let $\alpha^{(j)}$ be the line search step size, and let $\beta^{(j)}$ be an intercept with minimal train error. We perform gradient descent for each iteration $j \in \{0, 1, \dots\}$ via

$$\mathbf{w}^{(j+1)} = \mathbf{w}^{(j)} - \alpha^{(j)} \mathbf{X}^\top \bar{\mathbf{D}}_{\text{AUM}}(\mathbf{X}\mathbf{w}^{(j)} + \beta^{(j)}). \quad (30)$$

In the context of a more complex model such as a neural network, the AUM “gradient” $\bar{\mathbf{D}}_{\text{AUM}}$ can be used with back-propagation to obtain gradients with respect to all of the weights in the neural network. To regularize the model for the experiments in this paper, we use early stopping (number of iterations chosen by minimizing AUM or maximizing AUC with respect to a held-out validation set).

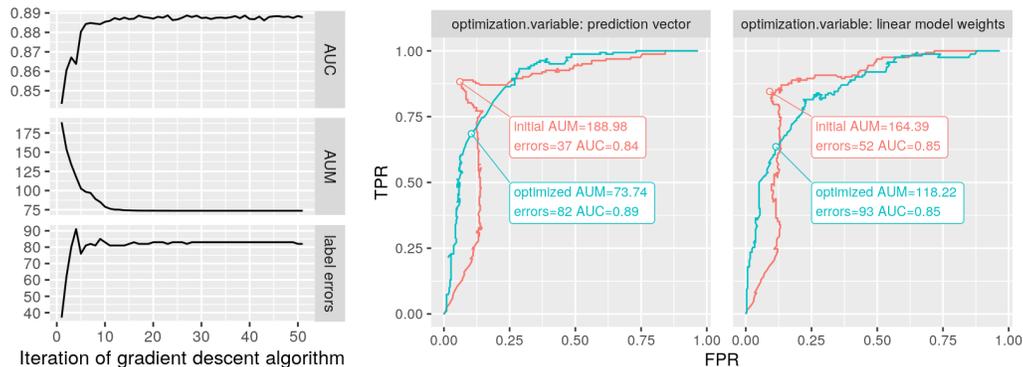


Figure 4: In a real changepoint detection problem with $n = 54$ labeled examples, minimizing the AUM tends to increase train AUC and label error rates. **Left:** we used AUM gradient descent on the n -vector of predicted values $\hat{\mathbf{y}}$ in the train set (initial predictions chosen by minimizing label errors for each example). Decreases in AUM happen during the same iterations as increases in AUC and label error rates. **Middle:** ROC curves before and after using AUM gradient descent on the n -vector of predicted values (dot shows a point on each ROC curve with minimum label errors). Although the AUM minimization resulted in higher label error rates, the ROC curve after the gradient descent appears much smoother (less sub-optimal points) than the initial ROC curve. **Right:** optimizing the p -vector of weights in a linear model ($p = 27$ features, initial weights minimize an un-regularized squared hinge loss summed over labeled examples). Note that AUC=0.85 is unchanged by the optimization but AUM decreases and error rate increases.

4. Empirical Results

We empirically study AUM minimization in the context of binary classification and changepoint detection problems. Our goal is to demonstrate that AUM minimization can result in AUC maximization with respect to train and test data. Note that for simplicity, we limit most of our comparisons to full gradient algorithms (each gradient computed using all training examples), although it is possible to use AUM with minibatch stochastic gradient algorithms (see Section 4.4).

4.1 AUM gradient descent increases train AUC and label errors in changepoint problems

In this experiment our goal was to demonstrate that minimizing the AUM results in maximizing the AUC in the train set. We used the chipseq data (a benchmark for labeled changepoint detection) from the UCI repository (Asuncion and Newman, 2007), treating each (set.name, fold) as a different train set, with pre-processing as previously described.¹ In brief, for each labeled example, changepoint models were computed for a range of penalty values, which resulted in models with a range of changepoints (some with few changepoints, others with many). Then the label error rate for each model was computed, along with a penalty $\hat{\lambda}_i$ which resulted in minimum label errors for each labeled example i . Finally these penalty values were used as the initial prediction vector $\hat{\mathbf{y}} = [-\log \hat{\lambda}_1 \cdots -\log \hat{\lambda}_n]$, which was used as the optimization variable in an AUM gradient descent algorithm. A line search was used for the step size so that the AUM was guaranteed to decrease at each iteration. Overall there were 68 different train sets with the number of labeled examples

1. <https://github.com/tdhock/feature-learning-benchmark>

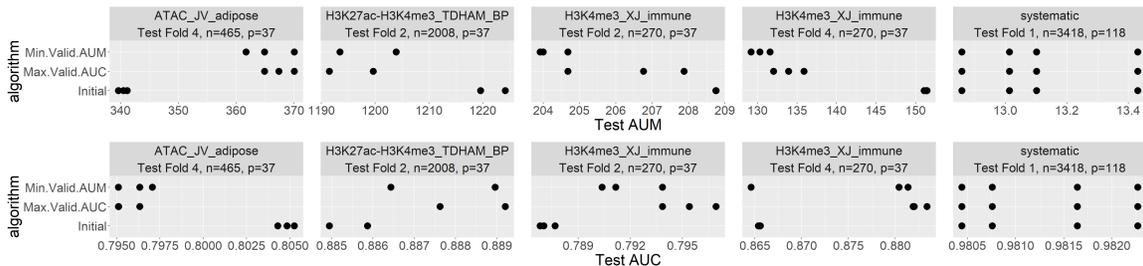


Figure 5: Test AUC is inversely correlated with test AUM values, when learning a linear model by AUM gradient descent. Algorithm (y axis) shows how number of iterations was chosen, by minimizing validation AUM, maximizing validation AUC, or taking the initial iteration. Each of five columns from left to right shows results for four random seeds on a given data set and test fold (n, p =number of examples and features for that entire data set).

ranging from $n = 7$ to 1011, and the total number of breakpoints in error functions ranging from $B = 47$ to 3104.

We expected that by using the predicted values directly as the optimization variable in gradient descent, we should be able to obtain ROC curves that were significantly different from the initialization (hopefully with larger AUC, even though the initialization came from minimizing the label error independently for each example). In a small number of train sets we observed that the optimization resulted in little or no change to both AUM and AUC; this happened when the initialization was close to or at a stationary point of the AUM (for example, when $AUM=0$ and $AUC=1$). However in most of the train sets we observed that minimizing the AUM results in increased AUC (on the train set).

In one representative train set (H3K4me3_XJ_immune fold 4 which has $n = 54$ labeled examples with a total of $B = 347$ breakpoints in error functions), we observed that the AUC and label error rate increases during the same iterations that the AUM decreases (Figure 4, left). Before optimization the ROC curve was highly non-monotonic with $AUC=0.84$ and a sharp point in the upper left corner; after AUM optimization the ROC curve became more regular with increased $AUC=0.89$ (Figure 4, middle). Interestingly, the AUC stayed below 1 after optimization, which suggests that AUM optimization does not result in looping ROC curves. This result suggests that AUM gradient descent can be used to maximize the AUC, although the label error rate also increases.

We performed a second experiment, this time optimizing $p = 27$ weights in a linear model parameter vector which was used to compute a prediction for each of the $n = 54$ labeled examples. The weights were initialized by using a gradient descent algorithm to minimize an un-regularized squared hinge loss that is a convex relaxation of the label error (Rigail et al., 2013). We expected the constraints of the linear model to reduce the accuracy with respect to the previous experiment (direct optimization of predicted values). We observed that after optimizing the weights using AUM gradient descent, the train AUC remained the same, but the train error rate increased (Figure 4, right). This experiment shows that the constraints of a linear model can prevent AUM minimization from resulting in AUC maximization (even in a data set for which it is possible to achieve larger AUC values).

4.2 Test AUM inversely correlated with test AUC in changepoint problems

In these experiments, the goal was to demonstrate that test AUC is inversely correlated with test AUM using our proposed linear model based on AUM minimization. We considered supervised

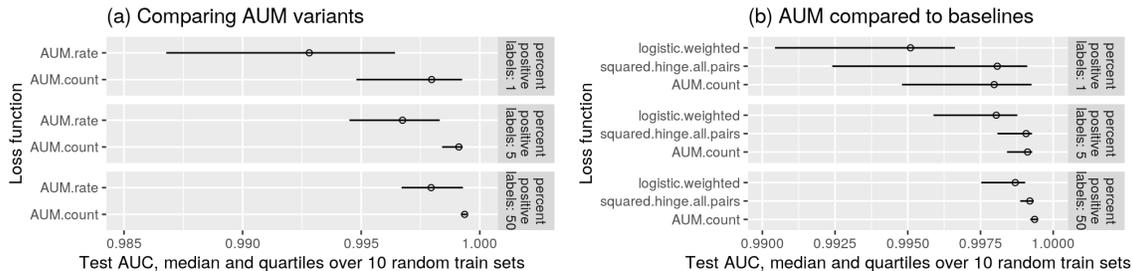


Figure 6: Test AUC in binary classification (zip data, images of 0/1 digits). Fixed test set with balanced classes (50% positive, 50% negative labels), and three train sets with different class imbalance (1%, 5%, 50% positive labels). **Left (a)**: AUM count is more accurate than AUM rate. **Right (b)**: AUM count minimization can result in poorer results with imbalanced datasets, but it is still at least as accurate as baselines, and sometimes more accurate.

changepoint detection data sets from a public repository.² It included pre-defined fold IDs that we used to define train/test splits over labeled examples of changepoint detection problems. In the five train sets that we analyzed (Figure 5), the number of labeled examples ranged from $n = 216$ to 3322, the number of features from $p = 26$ to 117, and the total number of error function breakpoints from $B = 1275$ to 7732. As explained in Section 3.5, our initialization and baseline was a linear model learned via gradient descent on a convex squared hinge loss with L1 regularization (Rigaill et al., 2013). To determine the extent to which the result depends on random initialization, we used four different seeds to select four different initial models. Some of the random seeds resulted in the same initial weights, despite having different seeds.

For each train set and random seed we used AUM gradient descent and selected the number of iterations via 4-fold cross-validation (take mean validation AUC/AUM over folds, then select iterations by maximizing AUC or minimizing AUM). In general we observed that both methods for selecting the number of iterations resulted in increased test AUC and decreased test AUM (middle three columns of Figure 5). One exception was for a single random initialization (fourth column, H3K4me3_XJ_immune, Min.Valid.AUM), both test AUC and test AUM decreased. For some data, the AUM gradient descent did not improve over the initialization (left and right columns). In one case there were very different train and test sets (left column, ATAC_JV_adipose Test Fold 4), so test AUM increased and test AUC decreased, consistent with our claim that AUM and AUC are inversely correlated. In another case (right column, systematic Test Fold 1), the selected number of gradient descent iterations was zero, so the test AUC/AUM of the learned model were the same as the initial model. In the cases where the AUC increased, the AUC values remained below 1, which provides evidence that AUM minimization does not create looping ROC curves. Overall these experiments in changepoint problems indicate that our proposed AUM gradient descent algorithm often results in increased test AUC, with respect to the initial/baseline linear model learned by minimizing a surrogate of the label error.

4.3 Test AUC in binary classification with linear model in zip data

The goal in this section was to study the accuracy of our proposed AUM loss function in unbalanced binary classification problems. Our experiment used the zip.train and zip.test image classification data³ from (Hastie et al., 2009). Each input is a 16×16 image which is represented by a vector

2. <https://github.com/tdhock/neuroblastoma-data>

3. <https://web.stanford.edu/~hastie/ElemStatLearn/>

$x \in [-1, 1]^{256}$. There are 10 classes (one for each digit), but we only used two classes (0/1) in order to study binary classification. For both the `zip.train` and `zip.test` files, we discarded some $y_i = 0$ labels in order to obtain sets with equal numbers of positive and negative labels (total 2010 in train, 528 in test). Then we generated 10 different train sets by randomly selecting 1000 examples with a class balance in $\{1\%, 5\%, 50\%\}$. The goal is to learn from the unbalanced train set data and then provide accurate predictions (as measured by AUC) on the balanced test set. The learning algorithms that we considered in our comparison were all based on gradient descent with constant step size and early stopping regularization. The step size and number of iterations hyper-parameters were chosen using grid search on a held-out validation set. For simplicity, we limit our comparisons in this section to other full gradient methods (not stochastic) which use other loss functions.

In our first comparison we wanted to understand if it is beneficial to normalize the AUM (use relative error rates rather than absolute error counts). In equations (11–12) we defined the AUM as the area under the minimum of false positive and false negative *counts*, but we could instead use *rates*. To compute AUM using rates, we need only change FPT/FNT functions in the minimum function (9) to FPR/FNR. In Figure 6 we refer to this variant as *AUM.rate*, and the original version as *AUM.count*. It is clear from Figure 6(a) that the *AUM.count* loss function variant has consistently larger test AUC than *AUM.rate*. The advantage of *AUM.count* becomes greater as the class imbalance is increased. For example, with no class imbalance median test AUC for *AUM.count* is 0.9993 and for *AUM.rate* is 0.9979 ($p = 0.03$ in one-sided t_9 -test); for large class imbalance with 1% positive labels median test AUC for *AUM.count* is 0.9979 and for *AUM.rate* is 0.9928 ($p = 0.004$). These results suggest that it is not beneficial to normalize the AUM, and it is more accurate to simply minimize the absolute $\min(\text{FP}, \text{FN})$ counts.

In our second comparison we wanted to compare AUM optimization to standard baseline loss functions for binary classification. As a baseline for margin loss functions (summed over labeled examples), we used a weighted version of the logistic loss $w_i \log[1 + \exp(y_i f(\mathbf{x}_i))]$. The weights $w_i \in \mathbb{R}$ are defined in order to ensure that the sum of weights over all negative/positive examples is the same (Menon et al., 2013). The total number of positive examples is $N_1 = |\{i : y_i = 1\}|$ and negative examples is $N_{-1} = |\{i : y_i = -1\}|$; the weights are therefore $w_i = 1/N_{y_i}$. For example when the class balance is 10%, there are $N_1 = 100$ positive examples and $N_{-1} = 900$ negative examples. Each positive example i has a weight of $w_i = 1/100$, and each negative example i has a weight of $w_i = 1/900$. The total weight over each class $y \in \{1, -1\}$ is $\sum_{i: y_i=y} w_i = 1$, which makes both classes equally important in the loss function. As a baseline for pairwise loss functions (summed over all pairs of positive $y_j = 1$ and negative $y_i = -1$ examples), we used a squared hinge loss, $[1 - f(\mathbf{x}_j) + f(\mathbf{x}_i)]_+^2$. Both baselines were chosen because they are supposed to maximize AUC (rather than the accuracy rate). We observed in Figure 6(b) that the learning algorithm using the AUM loss had test AUC values consistently better than or competitive with the baselines. The AUM loss had median test AUC that was consistently larger than the weighted logistic loss, with the largest difference for 1% positive labels (AUM=0.9979, logistic=0.9950, $p < 0.06$). There were smaller differences in median test AUC between the AUM loss and the squared hinge all pairs loss, with the largest difference for balanced labels (AUM=0.9993, pairs=0.9991, $p < 0.19$). Overall these experiments demonstrate that the AUM loss is highly competitive with the baseline loss functions, even when the train data set has large class imbalance.

4.4 Stochastic optimization of AUM for large binary classification problems using a convolutional neural network

In this section we show that our proposed loss functions can be used for learning using batched stochastic gradient descent in binary classification of large image data sets.

Implementation in PyTorch. We implemented the AUM loss functions (both count and rate) in PyTorch (Paszke et al., 2019), which supports automatic differentiation. We therefore only needed to implement the loss computation, and the gradient computation was provided automatically. Since

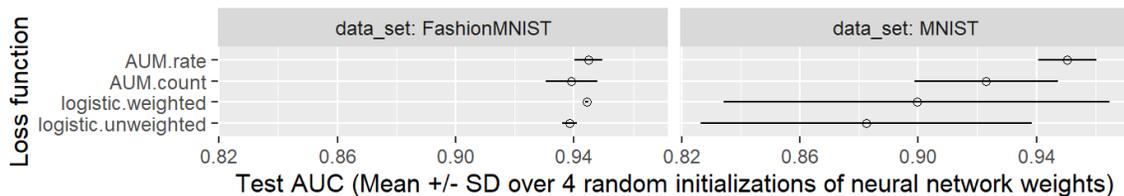


Figure 7: Test AUC in image classification experiments: test set was balanced (about 50% positive labels), whereas train set was imbalanced (about 1% positive labels).

the gradient is not defined everywhere for the AUM, PyTorch returns a subgradient in those cases. For example a simple non-differentiable point occurs for labels $y_1 = 0, y_2 = 1$ with predictions $\hat{y}_1 = \hat{y}_2 = 0$. In this case the directional derivatives are $[0, 1]$ for observation 1 and $[-1, 0]$ for observation 2, whereas PyTorch returned a derivative of -0 for each observation.

Data sets and algorithms used for comparison. We used MNIST and FashionMNIST data sets for these experiments (LeCun et al., 1998; Xiao et al., 2017). These data sets each have ten output classes to predict, so we created new binary labels by assigning label 0 to classes 0–4, and label 1 to classes 5–9. We used the test set of 10000 observations which was pre-defined, and we created train sets consisting of a subset of 1% of the positive examples (about 300), and all of the negative examples (about 30000). We then split each train set into 80% subtrain (about 24500 examples), and 20% validation (about 6100 examples). We used a version of the LeNet5 neural network architecture (LeCun et al., 1998), modified to use average pooling and ReLU activation functions.

Test AUC analysis. In this experiment we wanted to analyze the extent to which AUM minimization could result in effective learning under time constraints in large data. We therefore ran stochastic gradient descent for a maximum of 10 epochs, which means about 250 iterations/steps. We expected that AUM should result in highly accurate predictions as long as the batch size is large enough to allow one example of each class. We therefore used a batch size of 1000 (in order to ensure there is at least one positive example in each batch) and constant step sizes $\{10^{-4}, \dots, 10^2\}$. We choose the best number of iterations of gradient descent, and the step size, by maximizing AUC on the validation set. We observed that in both MNIST and FashionMNIST data sets, using the AUM loss resulted in test AUC values which were comparable to, or better than, the logistic loss (Figure 7). In the case of the MNIST data, the AUM.rate method shows a small but consistent improvement over the other methods (AUM.count mean test AUC difference = 0.027, p-value in one-sided t_3 test = 0.018; logistic.weighted difference = 0.051, p-value = 0.099; logistic.unweighted difference = 0.068, p-value = 0.061; logistic.unweighted means weight of $w_i = 1$ in the loss for each observation i). Overall these data provide convincing evidence that AUM can be used for highly accurate learning using neural networks in large data sets.

4.5 Speed comparison

The goal of this section is to show that AUM gradient computation has comparable speed to existing loss functions. In the case of binary classification, we compared the gradients of three loss functions, as shown in (Figure 8, left): squared hinge all pairs $O(n^2)$, AUM $O(n \log n)$, and logistic $O(n)$. For a fixed budget of computation time (1 second), the gradient of the AUM loss can be computed for much larger problems ($n = \text{millions}$) than the naïve square hinge all pairs approach ($n = \text{thousands}$). As expected, we observed that the AUM has a slightly slower asymptotic runtime (larger slope on log-log-plot) than the logistic loss.

We also did a speed test in the context of real changepoint detection problems, by comparing the AUM to a squared hinge loss summed over each of the n labeled examples (Rigaill et al., 2013). We

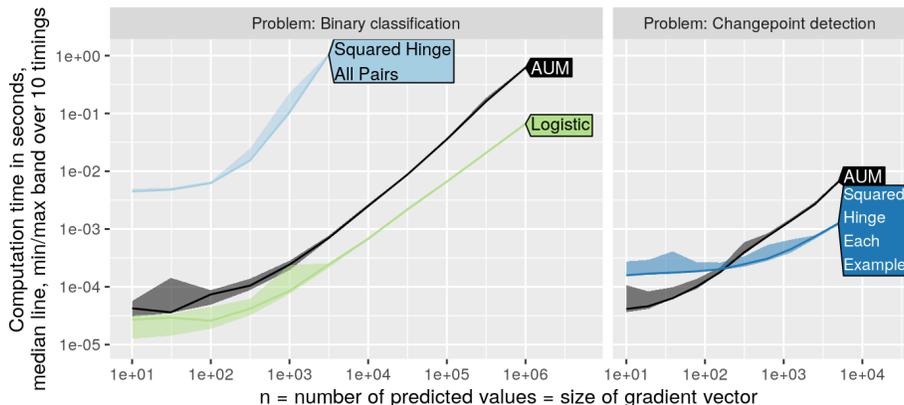


Figure 8: Speed comparison, time to compute gradient of various loss functions for data sets with variable number n of labeled examples. **Left:** for binary classification our timings suggest that AUM is $O(n \log n)$, faster than a $O(n^2)$ squared hinge loss summed over all pairs of positive/negative examples, and slower than $O(n)$ Logistic loss summed over all labeled examples. **Right:** for changepoint detection problems we observed that AUM is asymptotically similar to, but slower than, a squared hinge loss summed over labeled examples.

observed that the AUM gradient (log-linear time) is slightly slower than the squared hinge (linear time) in terms of asymptotic complexity (Figure 8, right). For example we observed about 10ms for AUM versus 1ms for the squared hinge for about $n \approx 5000$ examples with approximately 5 breakpoints in each example-specific error function. Overall these data show that the AUM has comparable speed to previously proposed loss functions for AUC optimization in binary classification, and changepoint detection.

5. Discussion and conclusions

In this paper we proposed the new AUM loss function which can be used in the context of prediction problems with false positive/negative rates such as in supervised binary classification and changepoint detection. We showed that the AUM can be interpreted as an L1 relaxation of a loss function that sums the $\min(\text{FP}, \text{FN})$ rate/count over all distinct points on the ROC curve. The AUM is highly novel with respect to previous loss functions, which either sum over all labeled examples (margin losses), or sum over all pairs of positive and negative examples (pairwise losses). Minimizing the AUM encourages the points on the ROC curve to be in the upper left of the (FPR, TPR) space. In section 2.8 we discussed an example of how $\text{AUM}=0$ (minimal) implies $\text{AUC}=1$ at differentiable points, which suggested using AUM as a surrogate loss for ROC curve optimization. We also discussed an example of how it is technically possible to have $\text{AUM}=0$ but AUC sub-optimal (at non-differentiable points). This degenerate case is not problematic in practice, because it is highly unlikely for the gradient descent learning algorithm to visit non-differentiable points.

We proposed a new algorithm for efficiently computing the AUM and its directional derivatives. For n labeled training examples and B total breakpoints in all error functions, our algorithm for computing the $n \times 2$ directional derivative matrix is $O(B \log B)$ time, which is much faster than previous $O(n^2)$ loss functions that naïvely sum over all pairs of positive and negative examples. In our empirical comparisons with several previous/baseline gradient descent learning algorithms, we observed comparable or better test AUC using our proposed method. We have shown that in large

imbalanced binary classification data sets, the AUM loss can be used with mini-batch stochastic gradient descent algorithms, and often out-performs baseline margin and pairwise losses. We have also observed that both the rate and count variants of AUM are promising surrogate loss functions; one or the other may result in more accurate learning, depending on the particular binary classification data set used. Overall for binary classification, our proposed AUM loss function gave comparable (if not better) results in terms of test AUC/AUM, with a much faster runtime than other algorithms that sum over all pairs of examples.

We observed that in changepoint detection problems with non-monotonic FP/FN functions, there is a tradeoff between AUC and accuracy that does not exist in binary classification problems. If maximizing accuracy with respect to the labels is important, we can use existing convex surrogates for label error such as the squared hinge loss (Rigaill et al., 2013). If maximizing AUC is important in labeled changepoint detection problems, then we suggest learning by minimizing our new AUM loss function which we have shown empirically results in AUC maximization (but lower accuracy). This tradeoff between AUC and accuracy means that the max accuracy model results in a highly non-monotonic ROC curve with many sub-optimal points, whereas the max AUC model has a more regular ROC curve (Figure 4). Overall our results showed that the AUM is a useful surrogate loss function for optimizing ROC curves in both binary classification and changepoint detection.

For future work, we have already started exploring the connection between AUM and other AUC relaxations, in order to create new AUC optimization algorithms. In fact, using an L1 relaxation pairwise loss (hinge with margin of zero) is quite similar to our approach (AUM integrates the min of FP and FN over all thresholds, whereas pairwise loss integrates the product). This suggests that it should be possible to compute a pairwise loss and gradient in log-linear time, so we are currently exploring a new algorithm based on that idea. Additionally, we would like to consider several variants of our new AUM loss function. It would be interesting to see if AUM could be effective for learning in the context of accelerated gradient descent algorithms such as SAGA (Gower et al., 2018), even though the AUM is not separable over training observations. Finally our current algorithm used either constant step size or a grid search, but a faster learning algorithm could potentially be obtained by exploiting the piecewise linear nature of the AUM during the step size computation.

Author statement. The authors declare that they have no conflict of interest. TDH contributed to the conception of ideas, proving theorems, implementing the algorithms, making figures, writing and revising the text. JH contributed to analysis of data, creation of figures, writing and revising the text.

Reproducible Research Statement. All of the software and data required to make the figures in this paper can be downloaded from <https://github.com/tdhock/max-generalized-auc>. An R package with C/C++ code that implements Algorithm 1 is available at <https://github.com/tdhock/aum>.

References

- Samaneh Aminikhanghahi and Diane J Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017.
- Arthur Asuncion and David Newman. UCI machine learning repository, 2007.
- Donald Bamber. The area above the ordinal dominance graph and the area below the receiver operating characteristic graph. *Journal of mathematical psychology*, 12(4):387–415, 1975.
- A Barnwal, H Cho, and TD Hocking. Survival regression with accelerated failure time model in xgboost. Preprint arXiv:2006.04920, 2020.
- Toon Calders and Szymon Jaroszewicz. Efficient AUC optimization for classification. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pages 42–53. Springer, 2007.

- Cristiano Leite Castro and Antonio Padua Braga. Optimization of the area under the ROC curve. In *2008 10th Brazilian Symposium on Neural Networks*, pages 141–146. IEEE, 2008.
- Corinna Cortes and Mehryar Mohri. AUC optimization vs. error rate minimization. *Advances in neural information processing systems*, 16(16):313–320, 2004.
- Alexandre Drouin, Toby Hocking, and Francois Laviolette. Maximum margin interval trees. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4947–4956. Curran Associates, Inc., 2017.
- James P Egan and James Pendleton Egan. *Signal detection theory and ROC-analysis*. Academic press, 1975.
- César Ferri, Peter Flach, and José Hernández-Orallo. Learning decision trees using the area under the ROC curve. In *ICML*, volume 2, pages 139–146, 2002.
- Atiyeh Fotoohinasab, Toby Hocking, and Fatemeh Afghah. A greedy graph search algorithm based on changepoint analysis for automatic qrs complex detection. *Computers in Biology and Medicine*, 130:104208, 2021. ISSN 0010-4825. doi: <https://doi.org/10.1016/j.combiomed.2021.104208>. URL <https://www.sciencedirect.com/science/article/pii/S0010482521000020>.
- Yoav Freund, Raj Iyer, Robert E Schapire, and Yoram Singer. An efficient boosting algorithm for combining preferences. *Journal of machine learning research*, 4(Nov):933–969, 2003.
- Robert Gower, Nicolas Le Roux, and Francis Bach. Tracking the gradients using the hessian: A new look at variance reducing stochastic methods. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 707–715. PMLR, 09–11 Apr 2018. URL <https://proceedings.mlr.press/v84/gower18a.html>.
- Guang Han and Chunxia Zhao. AUC maximization linear classifier based on active learning and its application. *Neurocomputing*, 73(7):1272 – 1280, 2010. ISSN 0925-2312. Advances in Computational Intelligence and Learning.
- David J Hand. Measuring classifier performance: a coherent alternative to the area under the roc curve. *Machine learning*, 77(1):103–123, 2009.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning*. Springer Series in Statistics. Springer, Springer Science+Business Media, LLC, 233 Spring Street, New York NY 10013, USA, second edition, 2009.
- Alan Herschtal and Bhavani Raskutti. Optimising area under the ROC curve using gradient descent. In *Proceedings of the twenty-first international conference on Machine learning*, page 49, 2004.
- Alan Herschtal, Bhavani Raskutti, and Peter K Campbell. Area under ROC optimisation using a ramp approximation. In *Proceedings of the 2006 SIAM International Conference on Data Mining*, pages 1–11. SIAM, 2006.
- TD Hocking and J Vargovich. Linear time dynamic programming for the exact path of optimal models selected from a finite set. arXiv:2003.02808, 2020.
- TD Hocking, V Boeva, G Rigaiil, G Schleiermacher, I Janoueix-Lerosey, O Delattre, W Richer, F Bourdeaut, M Suguro, M Seto, F Bach, and JP Vert. Seganndb: interactive web-based genomic segmentation. *Bioinformatics*, 30(11):1539–46, 2014.

- Toby Dylan Hocking and Guillaume Bourque. Machine Learning Algorithms for Simultaneous Supervised Detection of Peaks in Multiple Samples and Cell Types. In *Proc. Pacific Symposium on Biocomputing*, volume 25, pages 367–378, 2020.
- Toby Dylan Hocking, Gudrun Schleiermacher, Isabelle Janoueix-Lerosey, Valentina Boeva, Julie Cappo, Oliver Delattre, Francis Bach, and Jean-Philippe Vert. Learning smoothing models of copy number profiles using breakpoint annotations. *BMC Bioinformatics*, 14(164), May 2013.
- Toby Dylan Hocking, Guillem Rigai, and Guillaume Bourque. PeakSeg: constrained optimal segmentation and supervised penalty learning for peak detection in count data. In *Proc. 32nd ICML*, pages 324–332, 2015.
- Toby Dylan Hocking, Patricia Goerner-Potvin, Andreanne Morin, Xiaojian Shao, Tomi Pastinen, and Guillaume Bourque. Optimizing ChIP-seq peak detectors using visual labels and supervised machine learning. *Bioinformatics*, 33(4):491–499, 11 2016. ISSN 1367-4803.
- Toby Dylan Hocking, Guillem Rigai, Paul Fearnhead, and Guillaume Bourque. Constrained Dynamic Programming and Supervised Penalty Learning Algorithms for Peak Detection in Genomic Data. *Journal of Machine Learning Research*, 21(87):1–40, 2020.
- Sean W Jewell, Toby Dylan Hocking, Paul Fearnhead, and Daniela M Witten. Fast nonconvex deconvolution of calcium imaging data. *Biostatistics*, 02 2019. ISSN 1465-4644.
- Thorsten Joachims. A support vector method for multivariate performance measures. In *Proceedings of the 22nd international conference on Machine learning*, pages 377–384, 2005.
- Wojciech Kotlowski, Krzysztof Dembczynski, and Eyke Huellermeier. Bipartite ranking through minimization of univariate loss. In *ICML*, 2011.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Arnaud Liehrmann, Guillem Rigai, and Toby Dylan Hocking. Increased peak detection accuracy in over-dispersed ChIP-seq data with supervised segmentation models. *BMC Bioinformatics*, 22(323), 2021.
- Charles X Ling, Jin Huang, Harry Zhang, et al. Auc: a statistically consistent and more discriminating measure than accuracy. In *Ijcai*, volume 3, pages 519–524, 2003.
- Robert Maidstone, Toby Hocking, Guillem Rigai, and Paul Fearnhead. On optimal multiple changepoint algorithms for large data. *Statistics and Computing*, pages 1–15, 2016. ISSN 1573-1375.
- Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60, 1947.
- Pablo Martínez-Camblor, Norberto Corral, Corsino Rey, Julio Pascual, and Eva Cernuda-Morollón. Receiver operating characteristic curve generalization for non-monotone relationships. *Statistical methods in medical research*, 26(1):113–123, 2017.
- Aditya Menon, Harikrishna Narasimhan, Shivani Agarwal, and Sanjay Chawla. On the statistical consistency of algorithms for binary classification under class imbalance. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 603–611, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL <http://proceedings.mlr.press/v28/menon13a.html>.

- Harikrishna Narasimhan and Shivani Agarwal. A structural svm based approach for optimizing partial auc. In *International Conference on Machine Learning*, pages 516–524. PMLR, 2013.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Foster Provost and Tom Fawcett. Analysis and visualization of classifier performance with nonuniform class and cost distributions. In *Proceedings of AAAI-97 Workshop on AI Approaches to Fraud Detection & Risk Management*, pages 57–63, 1997.
- Alain Rakotomamonjy. Optimizing area under ROC curve with SVMs. In *ROCAI*, pages 71–80, 2004.
- Guillem Rigau, Toby Hocking, Jean-Philippe Vert, and Francis Bach. Learning sparse penalties for change-point detection using max margin interval regression. In *Proc. 30th ICML*, pages 172–180, 2013.
- R. Tyrrell Rockafellar. *Convex analysis*. Princeton Mathematical Series. Princeton University Press, Princeton, N. J., 1970.
- Cynthia Rudin, Corinna Cortes, Mehryar Mohri, and Robert E Schapire. Margin-based ranking meets boosting in the middle. In *International Conference on Computational Learning Theory*, pages 63–78. Springer, 2005.
- Clayton Scott. Calibrated asymmetric surrogate losses. *Electronic Journal of Statistics*, 6(none):958–992, 2012. doi: 10.1214/12-EJS699. URL <https://doi.org/10.1214/12-EJS699>.
- Charles Truong, Laurent Gudre, and Nicolas Vayatis. Penalty learning for changepoint detection. In *2017 25th European Signal Processing Conference (EUSIPCO)*, pages 1569–1573. IEEE, 2017.
- Gerrit JJ van den Burg and Christopher KI Williams. An evaluation of change point detection algorithms. *arXiv preprint arXiv:2003.06222*, 2020.
- Shijun Wang, Diana Li, Nicholas Petrick, Berkman Sahiner, Marius George Linguraru, and Ronald M Summers. Optimizing area under the roc curve using semi-supervised learning. *Pattern recognition*, 48(1):276–287, 2015.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. Preprint arXiv:1708.07747, 2017.
- Lian Yan, Robert H Dodier, Michael Mozer, and Richard H Wolniewicz. Optimizing classifier performance via an approximation to the Wilcoxon-Mann-Whitney statistic. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 848–855, 2003.
- Yiming Ying, Longyin Wen, and Siwei Lyu. Stochastic online auc maximization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 451–459, 2016.
- Zhuoning Yuan, Yan Yan, Milan Sonka, and Tianbao Yang. Robust deep auc maximization: A new surrogate loss and empirical studies on medical image classification. Preprint arXiv:2012.03173, 2020.

Peilin Zhao, Steven CH Hoi, Rong Jin, and Tianbo Yang. Online AUC maximization. In *28th international conference on machine learning*, 2011.